Inventory of Supporting Information

Manuscript #:
Corresponding author name: Susu Xu

Supplementary figure:

Figure #	Figure caption
Supplementary Fig. S1	Overview of monthly sales and repair for three storms.
Supplementary Fig. S2	Learned community networks for Hurricane Ian across representative communities stratified by socioeconomic and structural characteristics. Communities are grouped by income level (top row), parcel value (middle row), and community size (bottom row). Edge opacity reflects normalized influence intensity.
Supplementary Fig. S3	Learned community networks for Hurricane Ida across representative communities stratified by socioeconomic and structural characteristics. Communities are grouped by income level (top row), parcel value (middle row), and community size (bottom row). Edge opacity reflects normalized influence intensity.
Supplementary Fig. S4	Learned community networks for Hurricane Harvey across representative communities stratified by socioeconomic and structural characteristics. Communities are grouped by income level (top row), parcel value (middle row), and community size (bottom row). Edge opacity reflects normalized influence intensity.
Supplementary Fig. S5	Community characteristics that modulate neighbour-induced shifts across five post-disaster action pathways.
Supplementary Fig. S6	Sell and repair event state prediction accuracy on simulated data. Event-state prediction performance is evaluated using tail AUC on synthetic datasets generated from known diffusion graphs with controlled self- and neighbor-excitation dynamics. Results are reported across multiple simulation runs for both decision types and three simulation settings, using the same evaluation protocol as in the real-world case studies.
Supplementary Fig. S7	Sensitivity of forecasting performance to (a) self–excitation strength α_{self} and (b) neighbor–excitation strength α_{neigh} in the synthetic Hawkes simulator. Higher excitation levels produce stronger temporal or spatial signals, leading to substantially improved out-of-sample AUC for both sell and repair events.
Supplementary Fig. S8	Synthetic sensitivity analysis. (a) Increasing network connectivity k_{in} improves forecasting performance up to a moderate level. (b) Larger training ratios enhance AUC but exhibit diminishing returns beyond 75%.

Supplementary Table:

Table #	Table caption
Supplementary Table S1	Notation and metric definitions for various graph metrics utilized in the analysis.
Supplementary Table S2	Notation for household-level input characteristics.
Supplementary Table S3	Summary of benchmark model capabilities with respect to our three goals: future state prediction, network reconstruction, and mechanism (diffusion) recovery.
Supplementary Table S4	Diffusion-graph recovery on simulation data. Learned influence graphs are compared against the ground-truth diffusion graphs used to generate the simulated event sequences. Edge recovery is evaluated by (i) aligning simulated and learned communities, (ii) restricting comparison to parcels shared by both representations, and (iii) ensuring consistent edge orientation between the simulator and the learned adjacency. TP, FP, and FN denote true positive, false positive, and false negative directed edges, aggregated across evaluated communities and edge types.
Supplementary Table S5	Effect of graph connectivity k_{in} on event prediction and graph recovery performance.
Supplementary Table S6	Effect of self-excitation strength α_{self}.
Supplementary Table S7	Effect of neighbor excitation α_{neigh} on event prediction and graph recovery performance.
Supplementary Table S8	Effect of Hawkes temporal decay on event prediction and graph recovery performance.
Supplementary Table S9	Effect of diffusion horizon T on event prediction and graph recovery performance.
Supplementary Table S10	Effect of train–validation split ratio (TR) on event prediction and graph recovery performance.

Supplementary Discussion: Available below.

Supplementary Information

1 Supplementary Discussion

Overview of the Framework and Dataset

Fig. S1 provides more details of the monthly sale and repair data for each storm.

Influence Network Topology

Fig. S2 illustrates these contrasts for representative Hurricane Ian communities stratified by median household income, assessed parcel value, and improvement value. These structural differences complement the quantitative diffusion diagnostics and support the interpretation that heterogeneity in post-disaster diffusion is not only decision-specific (repairs versus sales) but also resource-stratified. In lower-resource settings, influence appears more distributed across many parcels, whereas in higher-resource settings it appears more centralized. Notably, these differences arise over similar spatial layouts, suggesting that resource stratification reshapes diffusion primarily through social and economic organization rather than geometry alone.

Fig. S3 visualizes representative parcel-level influence networks for Hurricane Ida communities stratified by income, parcel value, and community size. Consistent with the main-text results for Hurricane Ian, the inferred networks exhibit systematic heterogeneity in local connectivity and influence concentration across strata. Lower-income and lower-value communities tend to display denser and more locally redundant connectivity, with influence broadly distributed across short-range ties. Higher-income and higher-value communities exhibit sparser networks with visibly higher concentration of influence, where propagation is anchored around a limited number of salient parcels. Stratification by community size further reveals that smaller communities exhibit tightly knit local coupling, whereas larger communities show more directional and segmented influence pathways.

Fig. S4 presents analogous network visualizations for Hurricane Harvey, which experienced a more spatially extensive and prolonged recovery process. The inferred networks exhibit greater heterogeneity than those observed for Hurricane Ian or Ida, while preserving the same qualitative stratification by income,

parcel value, and community size. Lower-resource communities continue to show dense local connectivity with broadly shared influence, whereas higher-resource communities display sparser and more centralized influence structures. Compared to Ida, Harvey networks exhibit more elongated and heterogeneous influence pathways, consistent with the broader dispersion in diffusion strength observed in the supplementary diffusion diagnostics. Community size again modulates network organization, with larger communities exhibiting more directional and long-range influence pathways.

Simulation Experiments for Network Identifiability

To validate the proposed framework under controlled conditions, we conduct simulation experiments in which both the event-generation process and the underlying diffusion graph are known. Synthetic sell and repair sequences are generated from a graph-based Hawkes process with heterogeneous self- and neighbor-excitation, designed to reproduce the temporal burstiness and action-specific diffusion observed in post-disaster decision data. As shown in Fig. S6, our model achieves near-ceiling tail AUC across all simulated settings for both decision types, consistently outperforming spatial baselines, classical Hawkes variants, and neural temporal point-process models. This advantage is expected in simulation because the data-generating process contains structured neighbor-excitation on a latent graph; models that either ignore network structure or impose overly restrictive spatial kernels cannot fully exploit these dependencies, particularly in the distribution tail where rare-but-salient cascades occur.

Fig. S7 examines model sensitivity to the relative strength of self-excitation (α_{self}) and neighbor-excitation (α_{neigh}) in the synthetic Hawkes simulator. Increasing either excitation parameter systematically improves out-of-sample forecasting performance for both sell and repair decisions, indicating that the proposed framework effectively exploits stronger temporal or spatial dependency signals when they are present in the data. Self-excitation primarily governs temporal persistence at the parcel level. As α_{self} increases, events become more clustered in time, generating clearer autoregressive signatures that are readily captured by the model. Performance gains saturate at higher values, reflecting a regime in which future events become increasingly predictable from recent history alone. This saturation behavior suggests that once temporal clustering is sufficiently strong, additional self-excitation provides diminishing marginal information.

Neighbor-excitation controls cross-parcel influence and spatial propagation. When α_{neigh} is small, diffusion is weak and event occurrences are nearly independent across parcels, leading to poor predictability

and limited graph recovery. Once neighbor influence exceeds a moderate threshold, forecasting performance increases sharply, indicating that the model successfully identifies and leverages spatial contagion patterns. At very high α_{neigh}, performance again saturates, consistent with a diffusion regime dominated by strong network-driven cascades where uncertainty is primarily governed by initial seeds rather than stochastic propagation. Importantly, these trends hold consistently across both sell and repair event types, demonstrating that the framework does not rely on event-specific heuristics but instead responds appropriately to the underlying generative diffusion mechanism. The monotonic and saturating responses across excitation regimes confirm that the model is mechanism-sensitive rather than merely overfitting marginal event rates.

Fig. S8 evaluates sensitivity to graph connectivity (k_{in}) and the training data ratio. Increasing network connectivity initially improves forecasting accuracy by providing richer neighbor context and enabling the model to observe sufficient cross-parcel interactions to infer influence structure. However, performance gains taper at higher connectivity levels, reflecting a transition from informative local neighborhoods to increasingly redundant or noisy connections. This behavior indicates that the model benefits most from moderate, interpretable connectivity rather than dense, near-complete graphs. The training ratio sweep demonstrates that predictive performance improves steadily as more temporal data are observed, but with diminishing returns beyond moderate training fractions. This pattern suggests that the model rapidly learns the dominant diffusion structure and excitation dynamics from partial trajectories, after which additional data primarily refine parameter estimates without substantially altering the learned mechanism. The robustness to reduced training ratios highlights the data efficiency of the framework, which is particularly important in post-disaster settings where long time series may not be available. Taken together, these sensitivity analyses demonstrate that the proposed model responds coherently to changes in the true generative process. Performance improvements arise when diffusion signals are genuinely stronger or better supported by network structure and data availability, while saturation effects reveal principled limits rather than instability or overfitting. This controlled behavior supports the interpretability and reliability of the inferred diffusion mechanisms.

Table S3 contrasts representative modeling paradigms for post-disaster decisions along three objectives: future state prediction, network reconstruction, and diffusion mechanism recovery. Existing approaches typically address one or two of these goals, but not all simultaneously. Classical survival, competing-risks, and random-utility models are primarily designed for individual-level prediction [1,2]. While

they can incorporate covariates and neighborhood summaries, social influence enters only through aggregated or exogenously defined terms, precluding recovery of an explicit interaction structure or decomposition of self- versus neighbor-driven effects. Spatial econometric models parameterize spillovers through a fixed weights matrix, enabling estimation of average neighborhood influence strength, but the underlying network topology is imposed rather than inferred [3]. Agent-based recovery models encode rich behavioral rules and feedback mechanisms through simulation, capturing emergent dynamics but assuming, rather than learning, the diffusion process [4]. As a result, they do not support principled inference of latent influence pathways from observed decisions.

Beyond predictive accuracy, Table S4 provides a direct quantitative assessment of diffusion-graph reconstruction in a controlled simulation setting where the true influence structure is known. Learned influence networks are compared against the ground-truth diffusion graphs used to generate the event sequences, with evaluation restricted to consistently represented communities and households and with directed edge orientation matched between representations. Across all simulated hurricane scenarios, the recovered graphs achieve perfect precision (Precision = 1.000), indicating that the learned networks introduce no spurious directed influence ties, even when aggregated over more than 10^5 evaluated edges spanning 84–100 communities. This result demonstrates that the model is highly conservative in attributing social influence, reliably avoiding false-positive connections between households.

Tables S5–S10 provide a systematic sensitivity analysis that characterizes how event forecasting accuracy and diffusion-graph recovery respond to controlled perturbations of network structure, excitation balance, temporal memory, and data availability in a synthetic setting with known ground truth. Beyond validating robustness, these sweeps reveal fundamental distinctions between predictability and structural identifiability in diffusion-driven decision processes.

Varying the in-degree k_{in} (Table S5) exposes a clear transition from weakly informative to highly predictable diffusion regimes. When connectivity is minimal ($k_{\text{in}} = 1$), sell and repair prediction performance remains modest, particularly under windowed metrics, reflecting limited opportunity for neighbor-driven cascades. Once each node has even a small number of influential neighbors ($k_{\text{in}} \geq 3$), both tail-based and windowed AUC/AP rapidly increase and approach saturation, indicating that collective diffusion strongly enhances predictability.

In contrast, exact edge recovery becomes progressively more challenging as connectivity increases. Precision, recall, F1-score, and Jaccard index all decline with growing k_{in}, reflecting a fundamental identifiability constraint rather than model failure. In denser networks, multiple neighbors generate temporally correlated excitation patterns, so the model can reliably predict that neighborhood influence is present without being able to uniquely assign influence to a specific neighbor. As the number of true edges grows, recovering every minor interaction becomes statistically ill-posed, even though the dominant diffusion pathways remain predictive. This divergence highlights that high forecasting accuracy does not imply exhaustive edge-level recovery in dense diffusion environments.

Sweeping the self-excitation strength α_{self} (Table S6) yields steady improvements in both exact-timing and windowed prediction metrics for sell and repair events, consistent with stronger within-household temporal persistence. Structural recovery metrics remain effectively unchanged across this sweep, indicating that self-excitation primarily modulates intensity magnitude rather than altering cross-household influence attribution.

Neighbor excitation, by contrast, produces a pronounced regime shift (Table S7). When $\alpha_{\text{neigh}} \leq 1$, predictive performance is near chance, indicating dynamics dominated by exogenous variability and self-history. Once α_{neigh} exceeds this threshold, a cascade-dominated regime emerges, with rapid gains in both AUC and average precision. Notably, graph recovery metrics remain stable across this transition, implying that identifiability of influential neighbors depends more on topological sparsity and overlap than on the absolute magnitude of excitation. Together, these results disentangle the roles of self-memory and social coupling, showing that neighbor excitation governs whether diffusion is operative, while self-excitation governs how persistently decisions propagate over time.

Temporal decay (Table S8) and diffusion horizon T (Table S9) primarily regulate how influence accumulates over time. Moderate decay yields the best balance between exact and windowed prediction, whereas extremely slow decay accumulates long-past events and inflates exact AUC while destabilizing windowed metrics. Extending the observation horizon improves forecasting as cascades mature, with most gains realized by 36–60 months and diminishing returns thereafter. Across both sweeps, graph recovery metrics remain largely invariant, confirming that temporal parameters affect predictive signal aggregation rather than the underlying diffusion topology.

Varying the train, validation split ratio (Table S10) demonstrates robustness to reduced data availability. Even when only half of the temporal sequence is used for training, predictive performance remains high, and structural recovery metrics are stable across all splits. This indicates that influence networks become identifiable once a modest number of cascades has been observed, and that structural inference is constrained less by sample size than by the degree of confounding among candidate neighbors.

Taken together, these sensitivity analyses clarify how the proposed community-structured Graph Hawkes framework should be interpreted. Predictive accuracy is maximized in regimes with strong diffusion and sufficient temporal memory, whereas edge-level recovery is maximized when influence networks are sparse enough for neighbor identities to remain distinguishable. In dense or highly correlated environments, learned networks should be interpreted as capturing dominant diffusion pathways and regime-level influence rather than an exhaustive list of all minor ties. These results motivate reporting both forecasting and structural metrics, and they provide practical guidance for interpreting learned diffusion graphs in real-world disaster recovery settings.

2 Supplementary Figures

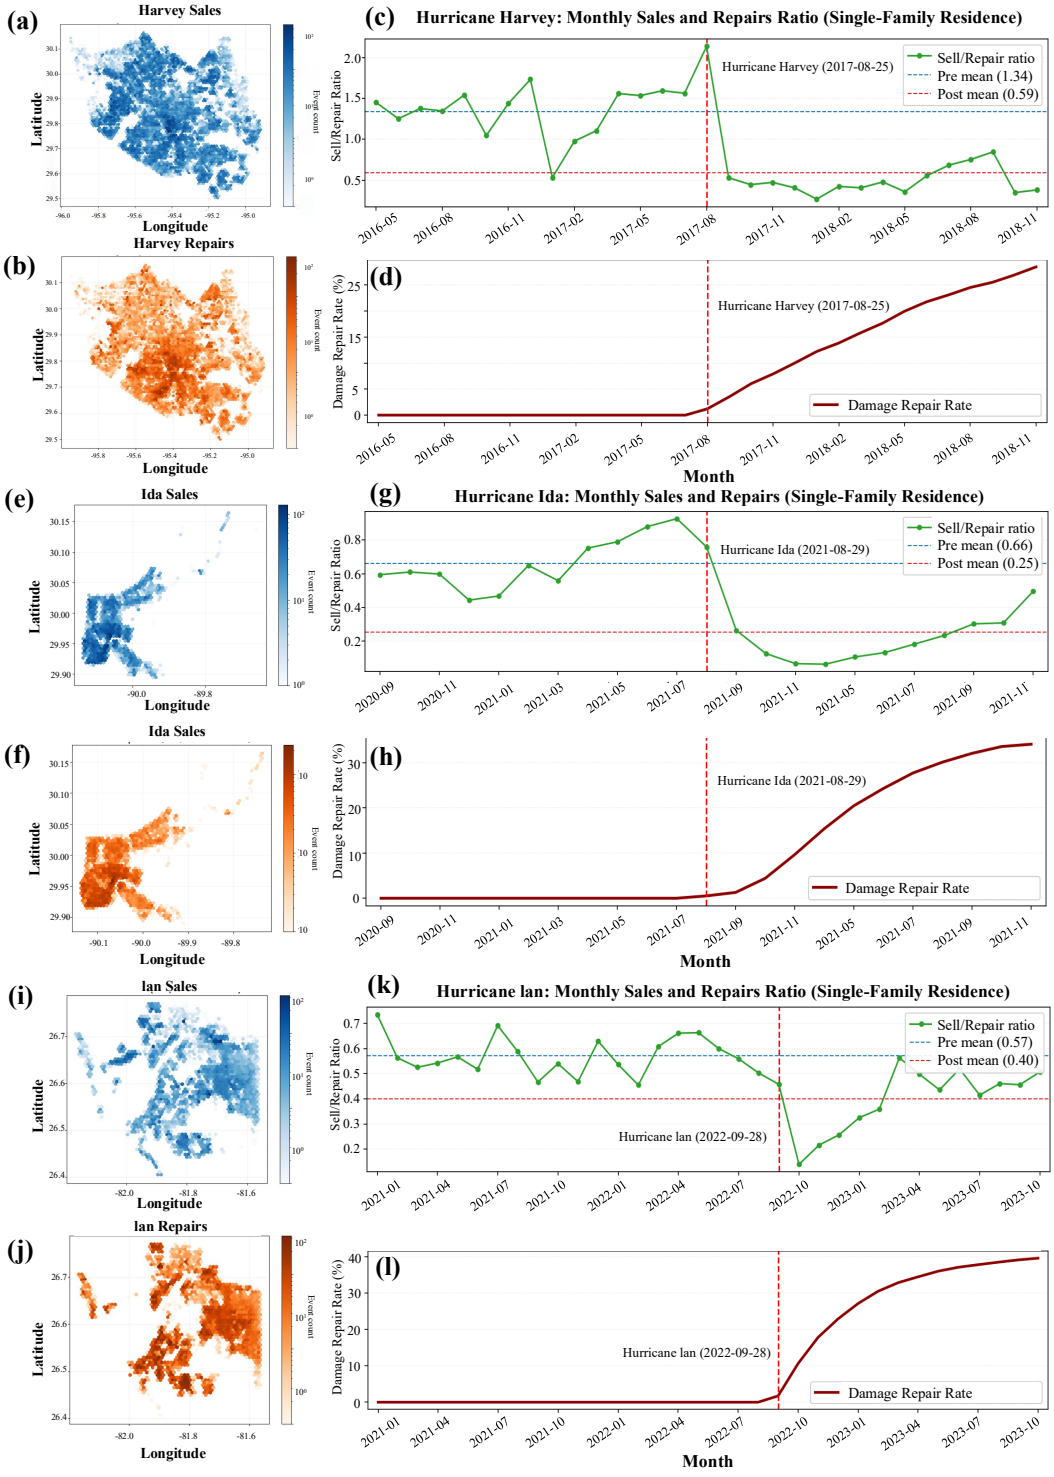


Fig. S1. Overview of sale and repair records. (a-d) Hurricane Harvey (Houston, TX; August 2017). (e-h) Hurricane Ian (Lee County, FL; September 2022). (i-l) Hurricane Ida (New Orleans, LA; August 2021).

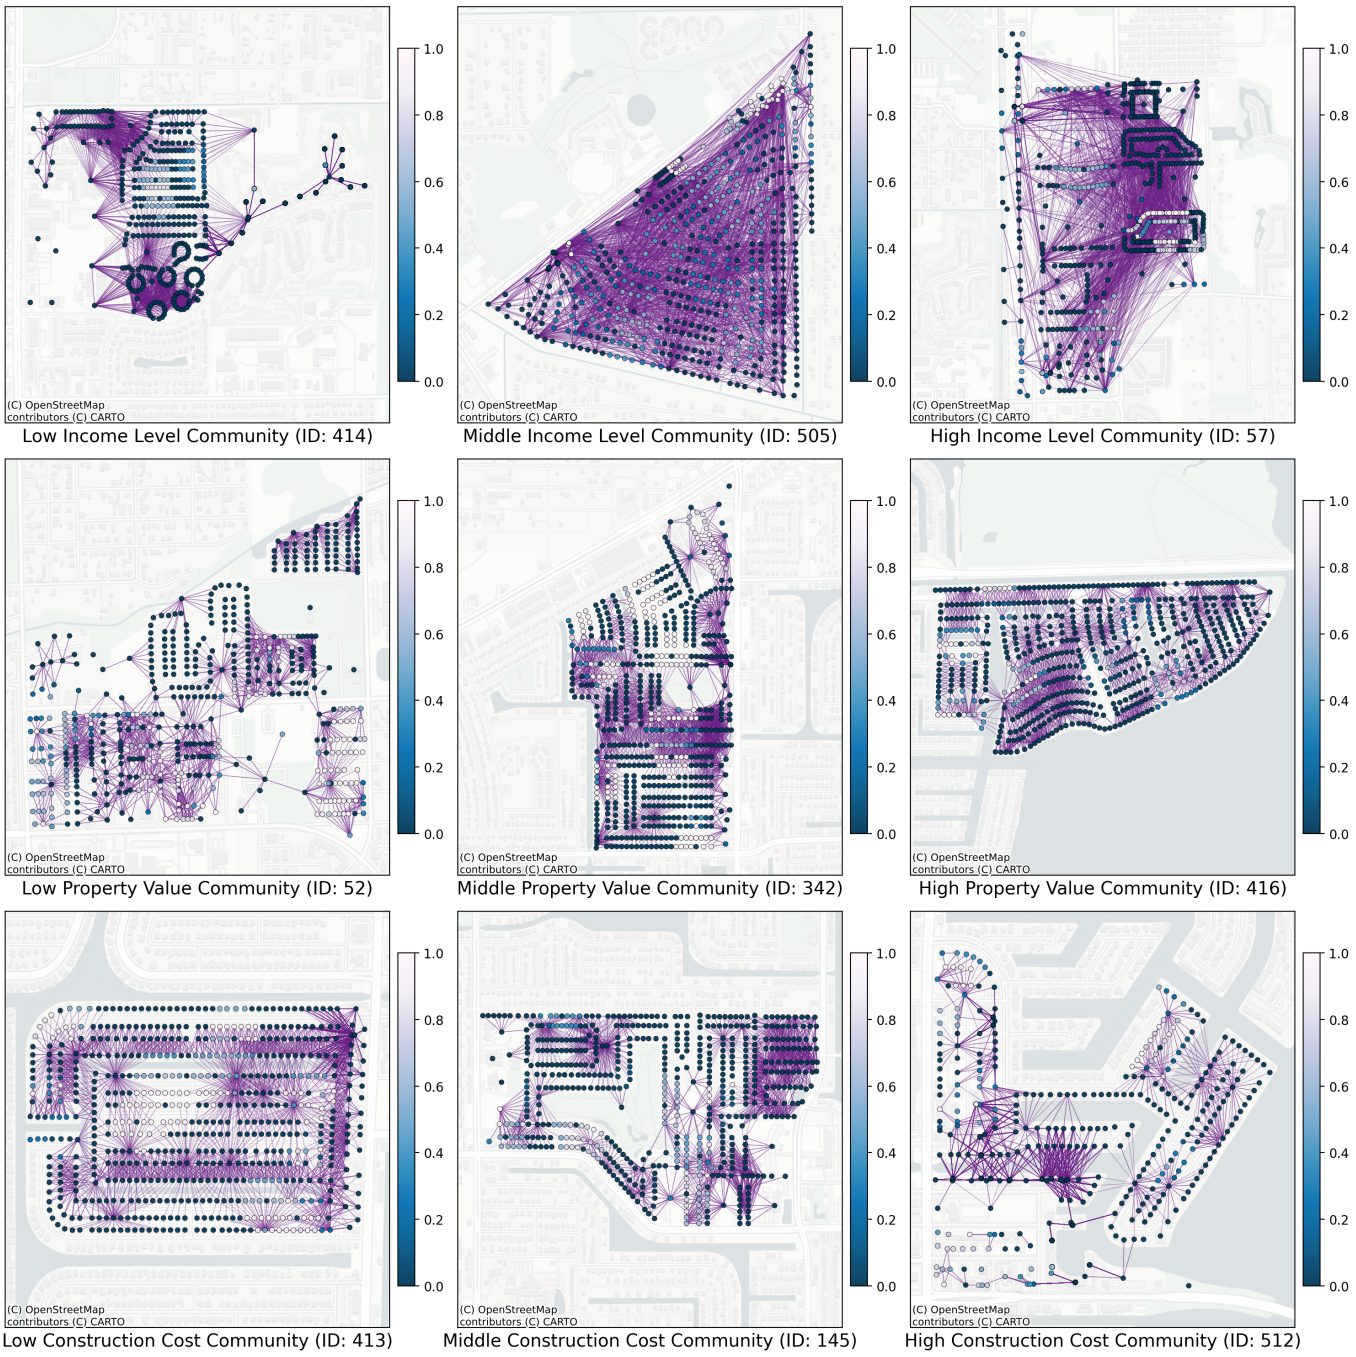


Fig. S2. Learned parcel-level influence networks for Hurricane Ian across representative communities. Panels visualize representative parcel-level diffusion graphs learned for Hurricane Ian communities stratified by income, property value, and construction cost. Nodes denote parcels and directed edges indicate inferred influence strength (opacity proportional to normalized weight). Readers should contrast the overall edge density and redundancy of local ties (left column) with the emergence of more hub-like, concentrated influence pathways in higher-resource communities (right column). These visual differences are quantified by edge-density and concentration statistics (out-strength inequality and top-edge mass).

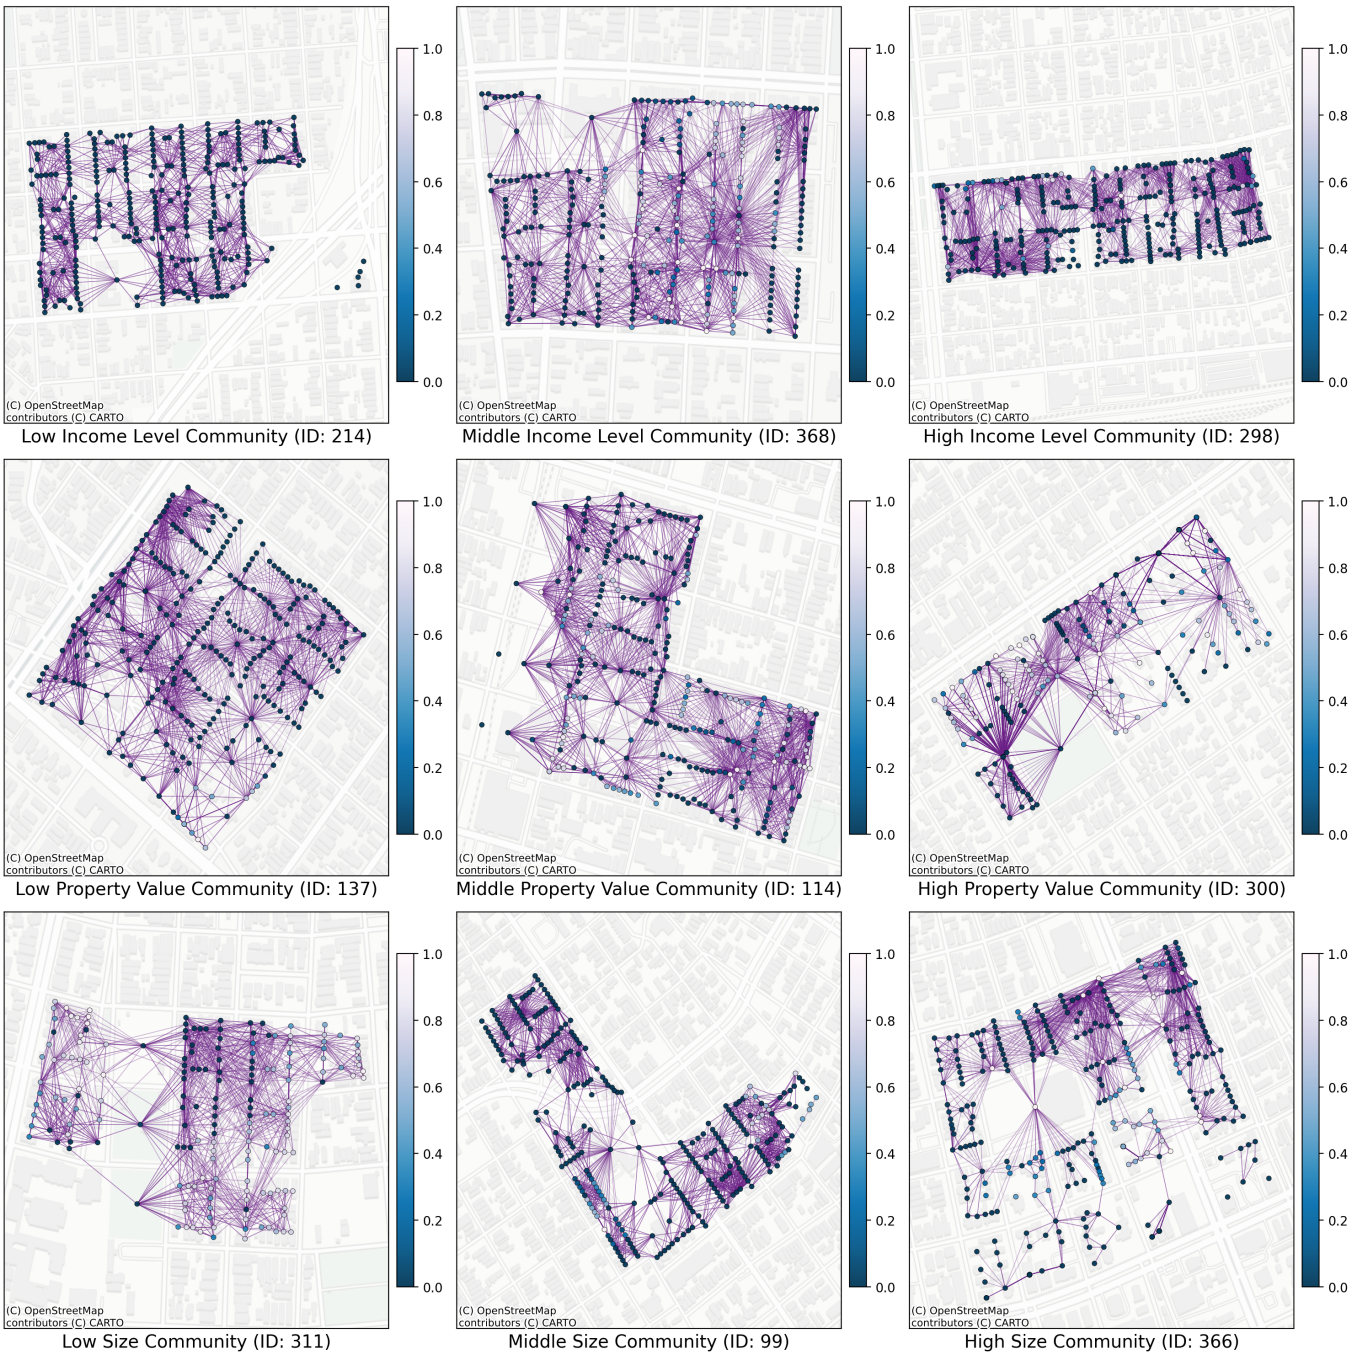


Fig. S3. Learned parcel-level influence networks for Hurricane Ida across representative communities. Each panel visualizes the inferred influence graph within a single community, where nodes represent residential parcels and directed edges denote learned social influence strength. Communities are stratified by median household income (top row), median parcel value (middle row), and community size (bottom row). Edge opacity reflects normalized influence intensity.

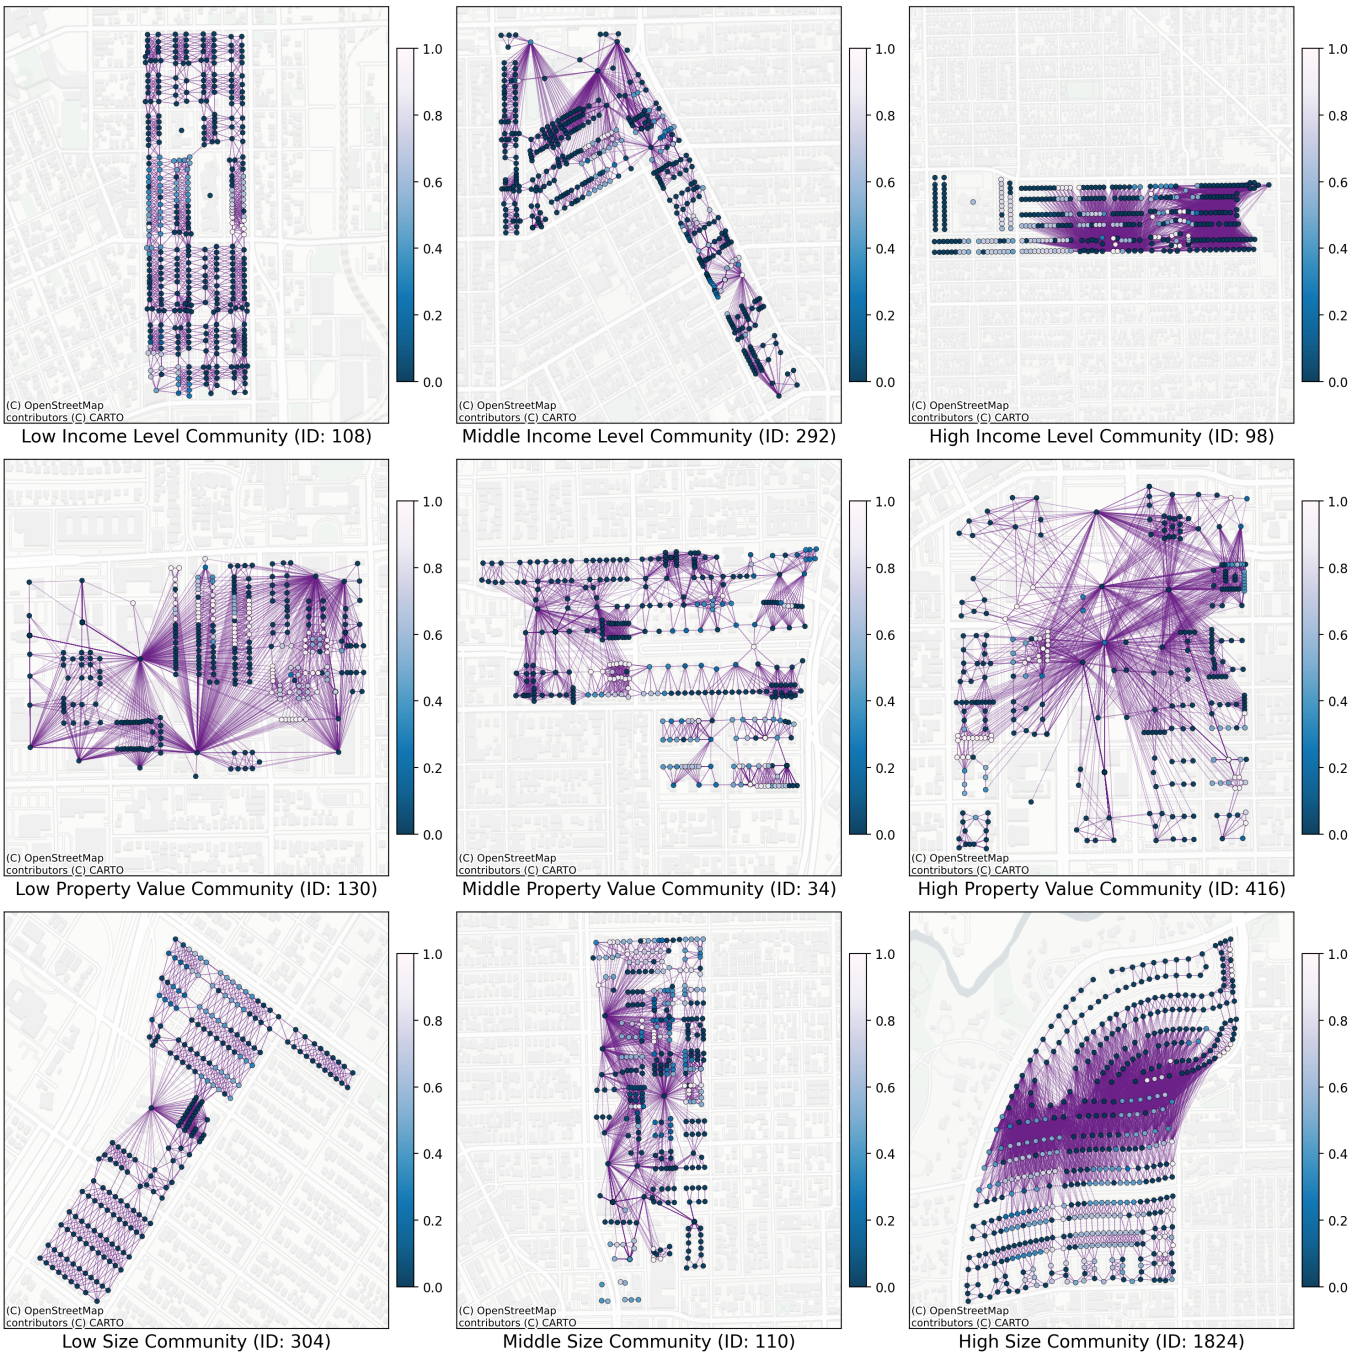


Fig. S4. Learned parcel-level influence networks for Hurricane Harvey across representative communities. Each panel visualizes the inferred influence graph within a single community, where nodes represent residential parcels and directed edges denote learned social influence strength. Communities are stratified by median household income (top row), median parcel value (middle row), and community size (bottom row). Edge opacity reflects normalized influence intensity.

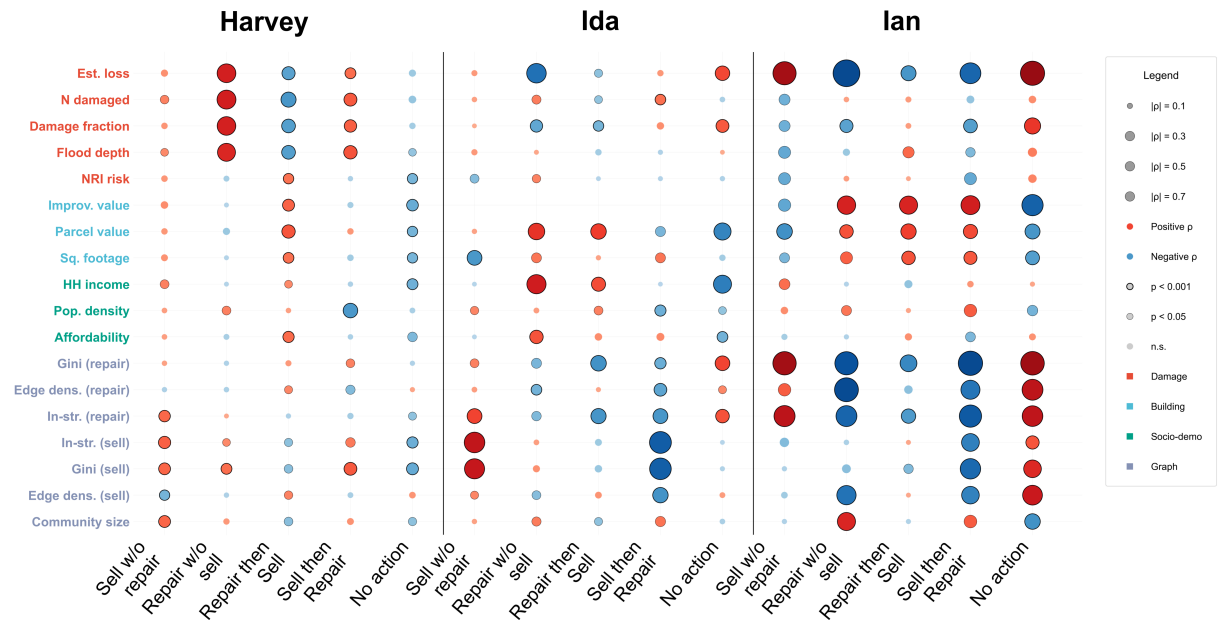


Fig. S5. Community characteristics that modulate neighbour-induced shifts across five post-disaster action pathways. Dot plot showing Spearman rank correlations between community-level features (rows) and the neighbour-induced probability shift for each action category (columns), separately for each hurricane. Dot size encodes correlation magnitude ($|\rho|$); colour encodes direction (red = positive, blue = negative). Black-edged dots indicate statistical significance (thick edge: $p < 0.001$; thin edge: $p < 0.05$). Feature labels are coloured by category: damage exposure (red), building characteristics (teal), socioeconomic and demographic factors (green), and graph structure of the learned neighbour network (grey-blue). In Harvey (left), damage exposure features (estimated loss, N damaged, damage fraction, flood depth) are the strongest predictors, primarily driving the repair-without-sell shift. In Ida (centre), sell-network topology (in-strength and Gini of the sell graph) predicts the sell-without-repair shift, while household income and parcel value predict resistance to no action. In Ian (right), repair-network structure (Gini, in-strength, edge density) dominates all five shifts, with correlations routinely exceeding $|\rho| = 0.5$. The shift from damage-dominated correlations in Harvey to graph-structure-dominated correlations in Ian indicates that when neighbour influence suppresses rather than activates recovery, the topology through which signals propagate becomes more important than the severity of physical damage.

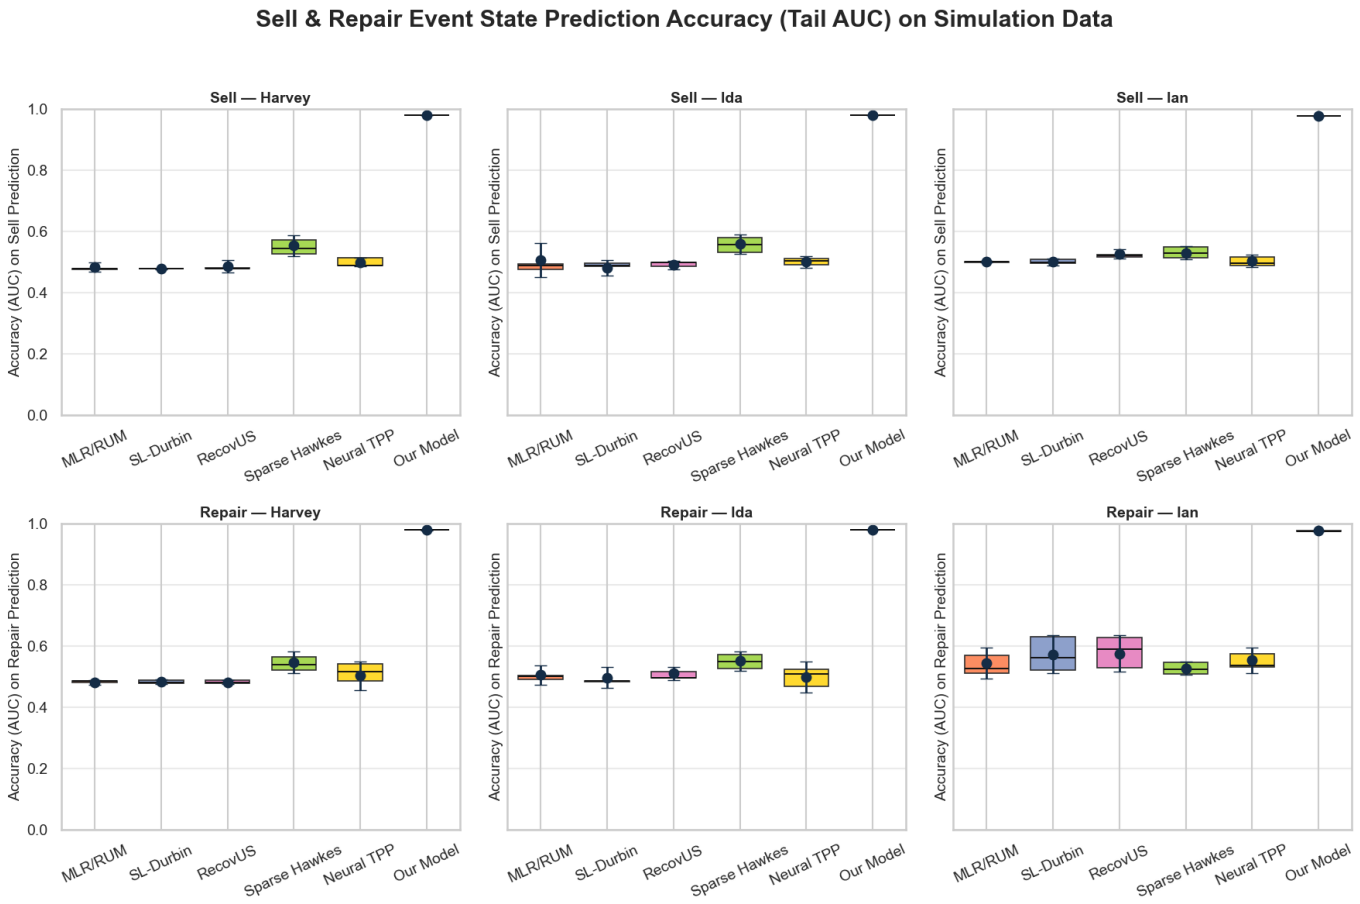


Fig. S6. Sell and repair event state prediction accuracy on simulated data. Event-state prediction performance is evaluated using tail AUC on synthetic datasets generated from known diffusion graphs with controlled self- and neighbor-excitation dynamics. Results are reported across multiple simulation runs for both decision types and three simulation settings, using the same evaluation protocol as in the real-world case studies.

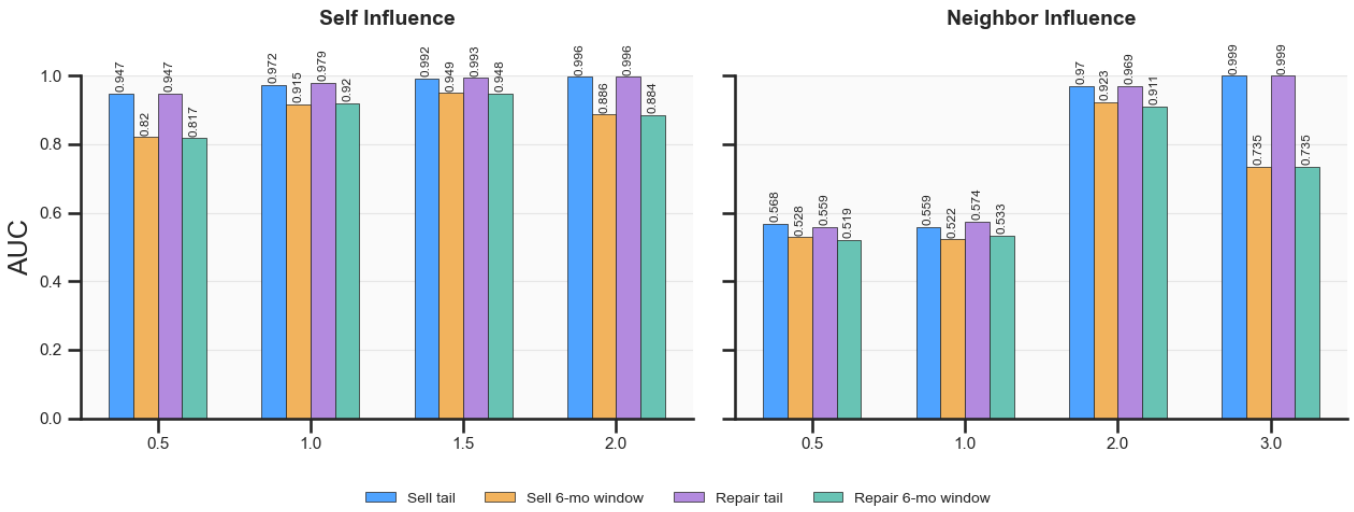


Fig. S7. Sensitivity of forecasting performance to (a) self–excitation strength α_{self} and (b) neighbor–excitation strength α_{neigh} in the synthetic Hawkes simulator. Higher excitation levels produce stronger temporal or spatial signals, leading to substantially improved out-of-sample AUC for both sell and repair events.

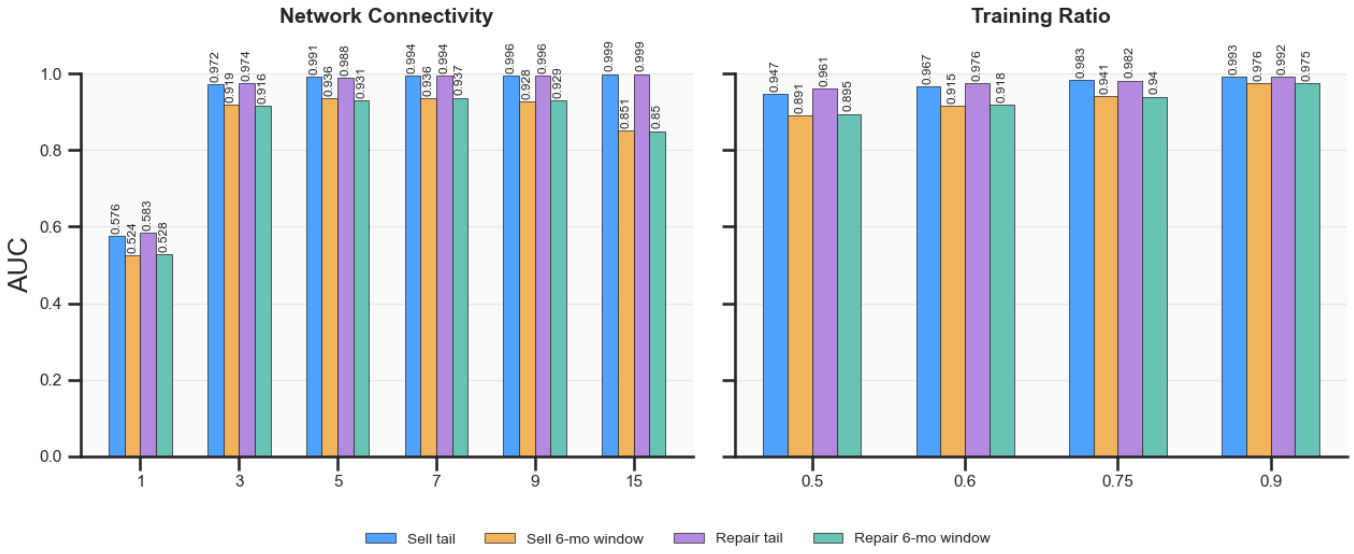


Fig. S8. Synthetic sensitivity analysis. (a) Increasing network connectivity k_{in} improves forecasting performance up to a moderate level. (b) Larger training ratios enhance AUC but exhibit diminishing returns beyond 75%.

3 Supplementary Tables

Supplementary Table S1. Notation and metric definitions. All graph metrics are computed per community on the learned adjacency matrix $\mathbf{A} \in \mathbb{R}^{2\times N\times N}$, where $\mathbf{A}_s = \mathbf{A}[0]$ (sell sub-graph) and $\mathbf{A}_r = \mathbf{A}[1]$ (repair sub-graph). N denotes the number of nodes (parcels) in a community.

Symbol	Name	Definition
Graph structure		
ρ_{edge}	Edge density	$\frac{\lvert\{(i,j) : A_{ij} > 0\}\rvert}{N^2}$; fraction of active pairwise connections.
ρ_{recip}	Reciprocity	$\frac{2\sum_{i,j}\min(A_{ij},\, A_{ji})}{\sum_{i,j} A_{ij}}$; fraction of total edge weight that is bidirectional.
f_{hub}	Hub fraction	$\frac{\sum_{i\in\mathcal{H}} s_i^{\text{out}}}{\sum_i s_i^{\text{out}}}$, where $\mathcal{H}$ is the top-10% of nodes by out-strength $s_i^{\text{out}} = \sum_j A_{ij}$.
C	Clustering coefficient	$\frac{1}{N}\sum_i \frac{t_i}{\binom{d_i}{2}}$, where t_i = triangles through node i and d_i = degree in the symmetrised binary graph $\mathbf{S} = \max(\mathbf{1}_{A>0},\, \mathbf{1}_{A>0}^\top)$. Computed as $t_i = \frac{1}{2}[\mathbf{S}^2 \circ \mathbf{S}]_{ii}$.
r_{mix}	Strength mixing	Pearson correlation $r(s_i^{\text{out}},\, s_j^{\text{out}})$ over all edges (i,j) with $A_{ij} > 0$; positive = assortative, negative = disassortative.
ϕ_{lwcc}	LWCC share	$\frac{\lvert\mathcal{C}_{\max}\rvert}{N}$; fraction of nodes in the largest weakly connected component.
$\bar{s}^{\text{out}}$	Mean out-strength	$\frac{1}{N}\sum_i s_i^{\text{out}} = \frac{1}{N}\sum_{i,j} A_{ij}$.
G	Gini coefficient	Gini index of the out-strength distribution $\{s_1^{\text{out}}, \ldots, s_N^{\text{out}}\}$; 0 = equal, 1 = maximal inequality.
λ_1	Spectral radius	Largest eigenvalue of $\mathbf{A}$; measures network amplification potential.
Cross-type structure		

Continued on next page

Symbol	Name	Definition
D_{sr}	Cross-type divergence	$\frac{\|\hat{\mathbf{A}}_s - \hat{\mathbf{A}}_r\|_F}{N}$, where $\hat{\mathbf{A}}_k = \mathbf{A}_k / \max(\mathbf{A}_k)$ is max-normalised.
α_{sr}	Cross-type alignment	Pearson $r(\mathbf{s}_s^{\text{out}}, \mathbf{s}_r^{\text{out}})$; correlation of sell and repair out-strength vectors across nodes.
Neighbor influence asymmetry		
$\Delta p(\text{S}\to\text{R})$	Sell-to-repair effect	Mean change in node i's repair probability when a neighbor's sell event is perturbed ($Y_j^s = 1$ vs. 0), averaged over nodes and time steps.
$\Delta p(\text{R}\to\text{S})$	Repair-to-sell effect	Mean change in node i's sell probability when a neighbor's repair event is perturbed.
$\text{Dom}_{\text{S}\to\text{R}}$	S→R dominance	$\frac{\|\{g : \|\Delta p_g(\text{S}\to\text{R})\| > \|\Delta p_g(\text{R}\to\text{S})\|\}\|}{G} \times 100\%$; share of communities where sell-to-repair channel dominates.
δ_{asym}	Asymmetry diff	$\|\Delta p(\text{S}\to\text{R})\| - \|\Delta p(\text{R}\to\text{S})\|$ per community; positive = S→R stronger.
Difference-in-differences (DiD) framework		
Δ_{aff}	Affected change	$\text{metric}_{\text{post, affected}} - \text{metric}_{\text{pre, affected}}$; within-group pre-to-post change for damaged communities.
Δ_{unaff}	Unaffected change	$\text{metric}_{\text{post, unaffected}} - \text{metric}_{\text{pre, unaffected}}$; counterfactual time trend.
τ_{DiD}	DiD estimand	$\Delta_{\text{aff}} - \Delta_{\text{unaff}}$; causal effect of hurricane damage net of common trends.

Supplementary Table S2. Household, property, neighborhood, and hazard-related characteristics used in the analysis.

Feature	Description
Property value	
improvval	Assessed value of the structure or building only.
landval	Assessed value of the land only.
parval	Total parcel assessed value, including land and structure.
agval	Agricultural value of the parcel, typically zero for residential properties.
tot appr val	total appraised value.
tot mkt val	total market value.
tot rcn val	replacement cost new value.
Building characteristics	
yearbuilt	Year the structure was built.
sqft	Building footprint area.
totalarea	Total lot or parcel area.
existing heatedarea	Conditioned interior floor area before the hurricane.
existing bathrooms	Number of bathrooms before the hurricane.
existing bedrooms	Number of bedrooms before the hurricane.
new heatedarea	Conditioned interior floor area after post-hurricane renovation or permitting activity.
new totalarea	Total parcel or building area after renovation or permitting activity.
new bathrooms	Number of bathrooms after renovation or permitting activity.
new bedrooms	Number of bedrooms after renovation or permitting activity.
story height	Building story height.
fireplaces	Number of fireplaces.
Location and neighborhood context	

Continued on next page

Feature	Description
lat, lon	Geographic coordinates of the parcel.
population density	Population density in the surrounding area.
median household income	Median household income of the census block group.
housing affordability index	Local housing affordability relative to income; higher values indicate greater affordability.
highest parcel elevation	Highest ground elevation within the parcel.
lowest parcel elevation	Lowest ground elevation within the parcel.
transmission line distance	Distance to the nearest power transmission line.
roughness rating	Terrain or vegetation roughness indicator relevant to wind and flood flow.
Growth projections	
population growth past 5 years	Percent population change over the previous 5 years.
population growth next 5 years	Projected percent population change over the next 5 years.
housing growth past 5 years	Percent change in housing units over the previous 5 years.
housing growth next 5 years	Projected percent change in housing units over the next 5 years.
household income growth next 5 years	Projected income growth over the next 5 years.
Risk classification	
fema nri risk rating	FEMA National Risk Index category, from very low to very high risk.

Continued on next page

Feature	Description
fema flood zone	FEMA flood zone designation, such as AE for 100-year floodplain or X for minimal flood risk.
Land use and ownership	
usecode	Local jurisdiction land-use code, such as single-family residential or commercial.
zoning	Zoning designation of the parcel.
lbcs activity	Land-Based Classification Standards code describing site activity.
lbcs function	Land-Based Classification Standards code describing parcel function.
lbcs structure	Land-Based Classification Standards code describing structural form.
owntype	Ownership type, such as individual, LLC, trust, or corporate owner.
Damage features	
FloodDepth	Estimated flood depth at the parcel.
EstLoss	Estimated monetary loss.
BldgValue	Building replacement value.
Occupancy	Occupancy type at the time of damage, such as single-family or condominium.
DamageLevel	FEMA categorical damage level: affected, minor, major, or destroyed.

Supplementary Table S3. Summary of benchmark model capabilities with respect to our three goals: future state prediction, network reconstruction, and mechanism (diffusion) recovery.

Model family	**Prediction**	**Network reconstruction**	**Mechanism recovery**
Cox / competing-risks survival	✓	×	∼ (aggregate neighbor effects)
Multinomial logit (RUM)	✓	×	∼ (aggregate neighbor effects)
Spatial lag / Durbin	✓	× (fixed W)	∼ (spillover strength only)
RecovUS-style ABM	✓(via simulation)	×	∼ (assumed rules, not inferred)
Sparse multivariate Hawkes	✓	✓	✓
Neural TPP (Neural Hawkes / RMTPP)	✓	×	∼ (implicit, black-box)
Our Graph Hawkes (community-structured)	✓	✓	✓

Supplementary Table S4. Diffusion-graph recovery on simulation data under matched top-k evaluation. Learned community-specific diffusion graphs are compared against the ground-truth graphs used to generate simulated event sequences. Evaluation aligns simulated and learned communities, restricts comparison to overlapping node sets, and matches edge orientation using `--transpose_pred`. To ensure a fair comparison with the simulator's in-degree constraint, predicted graphs are sparsified by keeping the top-k incoming edges per node with $k = k_{\text{in}}$ used during simulation. TP, FP, and FN denote directed edge counts aggregated across evaluated communities and decision types.

Setting	#Comm	k_{in}	top-k	TP	FP	FN	Precision	Recall	F1	Jaccard
Ian-sim	100	1	1	105291	0	105305	1.0000	0.5000	0.6666	0.5000
Ida-sim	84	1	1	55212	0	55224	1.0000	0.4999	0.6666	0.4999
Harvey-sim	100	1	1	91922	0	91922	1.0000	0.5000	0.6667	0.5000

Supplementary Table S5. Effect of graph connectivity k_{in} on event prediction and graph recovery performance.

k_{in}	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
1	0.576	0.083	0.583	0.089	0.524	0.023	0.528	0.024	1.000	0.500	0.667	0.500
3	0.972	0.951	0.974	0.950	0.919	0.798	0.916	0.797	0.337	0.253	0.289	0.169
5	0.991	0.986	0.988	0.982	0.936	0.850	0.931	0.844	0.208	0.173	0.189	0.104
7	0.994	0.993	0.994	0.993	0.936	0.882	0.937	0.884	0.155	0.136	0.145	0.078
9	0.996	0.996	0.996	0.996	0.928	0.894	0.929	0.894	0.127	0.114	0.120	0.064
15	0.999	1.000	0.999	1.000	0.851	0.937	0.850	0.937	0.095	0.089	0.092	0.048

Supplementary Table S6. Effect of self-excitation strength α_{self}.

α_{self}	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
0.5	0.947	0.898	0.947	0.894	0.820	0.618	0.817	0.611	0.337	0.253	0.289	0.169
1.0	0.972	0.949	0.979	0.956	0.915	0.795	0.920	0.800	0.337	0.252	0.289	0.169
1.5	0.992	0.990	0.993	0.990	0.949	0.903	0.948	0.902	0.337	0.253	0.289	0.169
2.0	0.996	0.999	0.996	0.999	0.886	0.936	0.884	0.935	0.337	0.252	0.289	0.169

Supplementary Table S7. Effect of neighbor excitation α_{neigh} on event prediction and graph recovery performance.

α_{neigh}	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
0.5	0.568	0.002	0.559	0.001	0.528	0.006	0.519	0.005	0.337	0.253	0.289	0.169
1.0	0.559	0.001	0.574	0.002	0.522	0.006	0.533	0.006	0.337	0.253	0.289	0.169
2.0	0.970	0.948	0.969	0.946	0.923	0.802	0.911	0.790	0.337	0.253	0.289	0.169
3.0	0.999	1.000	0.999	1.000	0.735	0.945	0.735	0.945	0.338	0.253	0.289	0.169

Supplementary Table S8. Effect of Hawkes temporal decay on event prediction and graph recovery performance.

Decay	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
0.3	0.943	0.889	0.946	0.891	0.822	0.616	0.820	0.613	0.337	0.253	0.289	0.169
0.5	0.972	0.948	0.975	0.954	0.923	0.804	0.923	0.803	0.337	0.253	0.289	0.169
0.7	0.994	0.995	0.995	0.996	0.944	0.918	0.943	0.917	0.337	0.252	0.289	0.169
0.9	0.999	1.000	0.999	1.000	0.638	0.945	0.638	0.945	0.338	0.253	0.289	0.169

Supplementary Table S9. Effect of diffusion horizon T on event prediction and graph recovery performance.

T	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
24	0.928	0.597	0.888	0.503	0.870	0.415	0.851	0.418	0.337	0.253	0.289	0.169
36	0.980	0.958	0.984	0.962	0.937	0.796	0.941	0.804	0.337	0.253	0.289	0.169
48	0.969	0.945	0.972	0.947	0.916	0.796	0.913	0.791	0.337	0.253	0.289	0.169
60	0.974	0.965	0.983	0.972	0.920	0.840	0.930	0.845	0.337	0.253	0.289	0.169

Supplementary Table S10. Effect of train–validation split ratio (TR) on event prediction and graph recovery performance.

TR	sAUC	sAP	rAUC	rAP	sAUCw	sAPw	rAUCw	rAPw	P	R	F1	J
0.50	0.947	0.920	0.961	0.934	0.891	0.758	0.895	0.759	0.338	0.253	0.289	0.169
0.60	0.967	0.944	0.976	0.952	0.915	0.797	0.918	0.793	0.337	0.252	0.289	0.169
0.75	0.983	0.969	0.982	0.967	0.941	0.812	0.940	0.805	0.337	0.253	0.289	0.169
0.90	0.993	0.980	0.992	0.978	0.976	0.750	0.975	0.749	0.337	0.253	0.289	0.169

References

1. Fine, J. P. & Gray, R. J. A proportional hazards model for the subdistribution of a competing risk. *Journal of the American statistical association* **94**, 496–509 (1999).

2. Train, K. E. *Discrete choice methods with simulation* (Cambridge university press, 2009).

3. LeSage, J. & Pace, R. K. *Introduction to spatial econometrics* (Chapman and Hall/CRC, 2009).

4. Moradi, S. & Nejat, A. RecovUS: An Agent-Based Model of Post-Disaster Household Recovery. *Journal of Artificial Societies and Social Simulation* **23**, 13 (2020).